



IPW

Docket No.: 566.43699X00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re the Application of:

Masayuki YAMAMOTO et al.

Serial No. 10/806,099

Filed: March 23, 2004

For: CONFIGURATION MANAGEMENT APPARATUS AND METHOD

REQUEST FOR RECONSIDERATION

August 30, 2005

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

In response to the Decision on Petition dated July 13, 2005,
reconsideration and withdrawal of the Decision is respectfully requested in view
of the following remarks.

REMARKS

Initially, in the Decision on Petition dated July 13, 2005, the Examiner notes that the Petition to Make Special filed June 6, 2005, and the Supplemental Petition filed on June 13, 2005, fail to submit a detailed discussion of the reference which points out how the claimed subject matter is patentable over the cited references.

It is submitted that the cited references, whether taken individually or in combination, fail to teach or suggest the invention as claimed. In particular, the cited references, at a minimum, fail to teach or suggest as recited in the claims:

a first feature of the present invention as recited in independent claim 1 wherein a third step for preparing a cache allocation definition of a logical volume of said second storage device which becomes a migration destination of a logical volume of said first storage device, in case that a cache allocation definition is included in the configuration obtained in said first step, on the basis of the cache allocation definition that said first storage device has and capacity of a cache that said second storage device has, and a fourth step for preparing a port bandwidth allocation definition of a port that said second storage device has, used for an access to a logical volume of said second storage device which becomes a migration destination of a logical volume of said first storage device, in case that a port bandwidth allocation definition is included in the configuration obtained in said first step, on the basis of the port bandwidth allocation definition that said first storage device has and bandwidth capacity of a port that said second storage device has;

a second feature of the present invention as recited in independent claim 5 including preparing a cache allocation definition of a logical volume of said second storage device which becomes a migration destination of a logical volume of said first storage device, in case that a cache allocation amount of a logical volume of said first storage device is smaller than remaining capacity to which a logical volume is not allocated among total capacity of the cache that said second storage device has, and prepares a port bandwidth allocation definition of a port that said second storage device has, which is used for an access to a logical volume of said second storage device which becomes a migration destination of a logical volume of said first storage device, in case that a port bandwidth allocation definition for an access to a logical volume of said first storage device is smaller than bandwidth remaining capacity which is not allocated to an access to a logical volume among total capacity of a port that said second storage device has;

a third feature of the present invention as recited in independent claim 6 wherein a process which prepares a cache allocation definition of a logical volume of said second storage device which becomes a migration destination of a logical volume of said first storage device, in case that a cache allocation definition is included in the configuration stored in said storage subsystem, on the basis of the cache allocation definition that said first storage device has and capacity of a cache that said second storage device has, and a process which prepares a port bandwidth allocation definition of a port that said second storage device has, which is used for an access to a logical volume of said second

storage device which becomes a migration destination of a logical volume of said first storage device, in case that a port bandwidth allocation definition is included in the configuration stored in said storage subsystem, on the basis of the port bandwidth allocation definition that said first storage device has and bandwidth capacity of a port that said second storage device has;

a fourth feature of the present invention as recited in independent claim 9 wherein a process which prepares a cache allocation definition of a logical volume of said second storage device which becomes a migration destination of a logical volume of said first storage device, in case that a cache allocation definition is included in the configuration which was stored in said storage subsystem, on the basis of the cache allocation definition that said first storage device has and capacity of a cache that said second storage device has, and a process which prepares a port bandwidth allocation definition of a port that said second storage device has, which is used for an access to a logical volume of said second storage device which becomes a migration destination of a logical volume of said first storage device, in case that a port bandwidth allocation definition is included in the configuration stored in said storage subsystem, on the basis of the port bandwidth allocation definition that said first storage device has and bandwidth capacity of a port that said second storage device has;

a fifth feature of the present invention as recited in independent claim 10 wherein a third step for preparing a cache allocation definition of a logical volume of said second storage device which becomes a migration destination of a logical volume of said first storage device, in case that a cache allocation definition is

included in the configuration obtained in said first step, on the basis of the cache allocation definition that said first storage device has and capacity of a cache that said second storage device has, and a fourth step for preparing a port bandwidth allocation definition of a port that said second storage device has, which is used for an access to a logical volume of said second storage device which becomes a migration destination of a logical volume of said first storage device, in case that a port bandwidth allocation definition is included in the configuration which was obtained in said first step, on the basis of the port bandwidth allocation definition that said first storage device has and bandwidth capacity of a port that said second storage has;

a sixth feature of the present invention as recited in independent claim 11 wherein a third step for preparing a cache allocation definition of a virtual volume of said virtualization device which virtualizes a logical volume of said first storage device, in case that a cache allocation definition is included in the configuration obtained in said first step, on the basis of the cache allocation definition that said first storage device has and capacity of a cache that said virtualization device has, and a fourth step for preparing a port bandwidth allocation definition of a port that said virtualization device has, which is used for an access to a virtual volume of said virtualization device which virtualizes a logical volume of said first storage device, in case that a port bandwidth allocation definition is included in the configuration which was obtained in said first step, on the basis of the port bandwidth allocation definition that said first storage device has and bandwidth capacity of a port that said virtualization device has;

a seventh feature of the present invention as recited in independent claim 16 wherein a process which prepares a cache allocation definition of a virtual volume of said virtualization device which virtualizes a logical volume of said first storage device, in case that a cache allocation definition is included in the configuration which was stored in said storage subsystem, on the basis of the cache allocation definition that said first storage device has and capacity of a cache that said virtualization device has, and a process which prepares a port bandwidth allocation definition of a port that said virtualization device has, which is used for an access to a virtual volume of said virtualization device which virtualizes a logical volume of said first storage device, in case that a port bandwidth allocation definition is included in the configuration stored in said storage subsystem, on the basis of the port bandwidth allocation definition that said first storage device has and bandwidth capacity of a port that said virtualization device has;

an eighth feature of the present invention as recited in independent claim 19 wherein a process which prepares a cache allocation definition of a virtual volume of said virtualization device which virtualizes a logical volume of said first storage device, in case that a cache allocation definition is included in the configuration stored in said storage subsystem, on the basis of the cache allocation definition that said first storage device has and capacity of a cache that said virtualization device has, and a process which prepares a port bandwidth allocation definition of a port that said virtualization device has, which is used for an access to a virtual volume of said virtualization device which virtualizes a

logical volume of said first storage device, in case that a port bandwidth allocation definition is included in the configuration which was stored in said storage subsystem, on the basis of the port bandwidth allocation definition that said first storage device has and bandwidth capacity of a port that said virtualization device has; and

a ninth feature of the present invention as recited in independent claim 20 wherein a third step for preparing a cache allocation definition of a virtual volume of said virtualization device which virtualizes a logical volume of said first storage device, in case that a cache allocation definition is included in the configuration obtained in said first step, on the basis of the cache allocation definition that said first storage device has and capacity of a cache that said virtualization device has, and a fourth step for preparing a port bandwidth allocation definition of a port that said virtualization device has, which is used for an access to a virtual volume of said virtualization device which virtualizes a logical volume of said first storage device, in case that a port bandwidth allocation definition is included in the configuration which was obtained in said first step, on the basis of the port bandwidth allocation definition that said first storage device has and bandwidth capacity of a port that said virtualization device has.

Further, the cited references fail to teach or suggest the above noted features of the present invention when taken in combination with other limitations recited in each specific claim (e.g., the first feature taken in combination with other limitations recited in claim 1, the second feature taken in combination with other limitations recited in claim 5, etc.).

The references considered most closely related to the claimed invention are briefly discussed below:

U.S. Patent No. 5,051,887 (Berger et al.) discloses a method and system for managing data storage devices and the data stored on them. The invention relates to improvements in data storage devices for the input and output of information to a data processing system and, more particularly, to a dual copy function in a disk storage device. Each write to the primary storage device is also written to a secondary storage device. Tables are employed to manage the records to be updated and their locations on the physical medium. A status array is also employed to retain the status and device identification information on a status track of each of the data storage devices and another location to provide global identification and management of interchangeable data storage devices. The method comprises means for selecting two direct access storage devices as duplex-paired storage devices, a first one of the direct access storage devices being designated as a primary storage device and a second one of the direct access storage devices being designated as a secondary storage device, the primary and secondary storage devices being physically separated but substantially identical in configuration and type. (See, e.g., Abstract and column 2, lines 20-50.) However, Berger et al. at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 5, the above described third feature of

the present invention as recited in independent claim 6, the above described fourth feature of the present invention as recited in independent claim 9, the above described fifth feature of the present invention as recited in independent claim 10, the above described sixth feature of the present invention as recited in independent claim 11, the above described seventh feature of the present invention as recited in independent claim 16, the above described eighth feature of the present invention as recited in independent claim 19 and the above described ninth feature of the present invention as recited in independent claim 20.

U.S. Patent No. 5,784,703 (Muraoka et al.) discloses a storage array system having a plurality of physical storage devices from and into which data are read and written by a host system, comprising a physical system configuration information table which stores configuration information of the plurality of physical storage devices. The system includes a logical system configuration information table which stores configuration information of a plurality of logical storage subsystems in the case where the plurality of physical storage devices have been divided into the logical storage subsystems; and a division designation device which effects at least one of a number of designations in compliance with a command received from the host system. The designations include an n-equal-part designation in which a total storage capacity of the physical storage devices is equally divided in n (an integer of at least 2) for the n logical storage subsystems, a divisional allotment rate designation in which the total storage capacity is divided in accordance with a specified number for

division and specified rates of capacities to be allotted to the respective logical storage subsystems, and a physical boundary division designation in which the total storage capacity is divided in accordance with positions for the division specified in physical storage unit. (See, e.g., Abstract and column 1, line 45, through column 2, line 49.) However, unlike the present invention, Muraoka et al., at a minimum, does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 5, the above described third feature of the present invention as recited in independent claim 6, the above described fourth feature of the present invention as recited in independent claim 9, the above described fifth feature of the present invention as recited in independent claim 10, the above described sixth feature of the present invention as recited in independent claim 11, the above described seventh feature of the present invention as recited in independent claim 16, the above described eighth feature of the present invention as recited in independent claim 19 and the above described ninth feature of the present invention as recited in independent claim 20.

U.S. Patent No. 6,289,423(Ozaki et al.) discloses a method and apparatus to control copying of a group of data stored in an original information-storing medium into a destination information-storing medium. The method includes a step of reading attribute information of a first zone storing the group of data, where the attribute information indicates whether the zone is a ROM area or a RAM area, and a step of copying the attribute information to a second zone

to which the group of data is to be copied. The method further includes a step of copying the group of data stored in the first zone in the original information-storing medium to the second zone in the destination information-storing medium. (See, e.g., Abstract and column 2, line 62, through column 3, line 7.) However, unlike the present invention, Ozaki et al., at a minimum, does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 5, the above described third feature of the present invention as recited in independent claim 6, the above described fourth feature of the present invention as recited in independent claim 9, the above described fifth feature of the present invention as recited in independent claim 10, the above described sixth feature of the present invention as recited in independent claim 11, the above described seventh feature of the present invention as recited in independent claim 16, the above described eighth feature of the present invention as recited in independent claim 19 and the above described ninth feature of the present invention as recited in independent claim 20.

U.S. Patent No. 6,775,739 (Bachmat et al.) discloses a mass storage system in which stored logical volumes are duplicated in mirrored form. The system includes a method for dynamically maximizing separation of the logical volumes by adjusting the mirror service policy for a disk drive system. Statistics are periodically collected describing the reading and writing of data to mirrored logical volumes of the system in successive time periods and, from time to time,

based upon the collected statistics, activity levels for the logical volumes are determined. A matching algorithm, weighted or unweighted, maximizes separation of the logical volumes and minimizes average physical drive seek times. In particular, the system takes into account activity levels at the physical devices and results in more efficient accessing of logical volume pairs as well as a better balance or loading of logical volumes. (See, e.g., Abstract and column 1, line 63, through column 2, line 58.) However, unlike the present invention, Bachmat et al., at a minimum, does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 5, the above described third feature of the present invention as recited in independent claim 6, the above described fourth feature of the present invention as recited in independent claim 9, the above described fifth feature of the present invention as recited in independent claim 10, the above described sixth feature of the present invention as recited in independent claim 11, the above described seventh feature of the present invention as recited in independent claim 16, the above described eighth feature of the present invention as recited in independent claim 19 and the above described ninth feature of the present invention as recited in independent claim 20.

U.S. Patent Publication No. 2002/0138705 (Suzuki et al.) discloses a data saving method, apparatus and computer program which operates in an information processing system including a first external storage device, a host device which performs data input and/or output operations to and/or from the first

external storage device and a second external storage device. A connection is formed between the first external storage device and the second external storage device. The second external storage device is not connected to the host device. An operation is performed for automatically saving data from the first external storage into the second external storage device via the connection formed between the first external storage device and the second external storage device. The saving is performed based on configuration information stored in the first external storage device. The configuration information includes information for designating an execution start time of saving data to the second external storage device, and at least one of a data storing area and a data set name in the second external storage device at which stored data is to be stored and information for indicating at least one type of second external storage device. (See, e.g., Abstract and paragraphs 5-7.) However, unlike the present invention, Suzuki et al., at a minimum, does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 5, the above described third feature of the present invention as recited in independent claim 6, the above described fourth feature of the present invention as recited in independent claim 9, the above described fifth feature of the present invention as recited in independent claim 10, the above described sixth feature of the present invention as recited in independent claim 11, the above described seventh feature of the present invention as recited in independent claim 16, the above described eighth feature of the present invention as recited in independent claim

19 and the above described ninth feature of the present invention as recited in independent claim 20.

U.S. Patent Publication No. 2003/0221063 (Eguchi et al.) discloses a storage subsystem that detects the necessity of the data relocation and determines whether the data relocation is possible or not from one storage unit to another within the particular storage subsystem. A virtualization device is notified in the case where the data relocation in the particular subsystem is impossible. The manager of the virtualization device gives an inquiry to the storage subsystems as to whether any one of them can become a relocation destination or not. A storage subsystem determines, based on the related information, whether the data relocation requested is possible or not within the particular storage subsystem. In the case where such relocation is possible, the copying process is carried out for data relocation from one storage subsystem to another in compliance with an instruction from the virtualization device. The data relocation method for a computer system including a host computer, a plurality of storage subsystems each having a plurality of storage units, and a data relocation management unit for managing the data relocation between the storage subsystems, comprises the steps of: determining whether the data relocation from one storage unit to another within one of the storage subsystems which constitutes a relocation source, based on the performance information and the utilization information of the particular one of the storage subsystems; determining whether the data relocation is possible or not to a storage unit within another one of the storage subsystems through the data relocation management

unit in the case where the data relocation is impossible within the storage subsystem constituting the relocation source; and copying the data to the particular one storage subsystem constituting a relocation destination from the storage subsystem constituting the relocation source through the data relocation management unit in the case where the data relocation to the particular another storage subsystem is possible. (See, e.g., Abstract and paragraphs 18-21.) However, unlike the present invention, Eguchi et al., at a minimum, does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 5, the above described third feature of the present invention as recited in independent claim 6, the above described fourth feature of the present invention as recited in independent claim 9, the above described fifth feature of the present invention as recited in independent claim 10, the above described sixth feature of the present invention as recited in independent claim 11, the above described seventh feature of the present invention as recited in independent claim 16, the above described eighth feature of the present invention as recited in independent claim 19 and the above described ninth feature of the present invention as recited in independent claim 20.

U.S. Patent Publication No. 2004/0024977 (Delaney et al.) discloses a method and system for copying and/or transferring stored data of one storage volume of a storage system to another storage volume while enabling requests to the storage volumes. The system and method is particularly useful in RAID

storage systems. One system includes a detector configured for detecting an amount of the stored data copied from the one storage volume to the other storage volume based on a boundary marker of the storage volumes. The boundary marker may indicate an amount of the stored data copied from the one storage volume to the other storage volume. A processor is communicatively connected to the detector for processing requests to the storage volumes according to rules based on the boundary marker. A file generator may generate a log file of the requests. The log file may be stored in the other storage volume thereby minimizing a need for external storage devices. (See, e.g., Abstract and paragraphs 12-15.) However, unlike the present invention, Delaney et al., at a minimum, does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 5, the above described third feature of the present invention as recited in independent claim 6, the above described fourth feature of the present invention as recited in independent claim 9, the above described fifth feature of the present invention as recited in independent claim 10, the above described sixth feature of the present invention as recited in independent claim 11, the above described seventh feature of the present invention as recited in independent claim 16, the above described eighth feature of the present invention as recited in independent claim 19 and the above described ninth feature of the present invention as recited in independent claim 20.

U.S. Published Patent Application 2004/0123029 to Dalal et al.,

discloses a method and system for preserving the intent of a logical volume creator with the logical volume. The creator of a logical volume defines certain characteristics of that logical volume, such as a number of copies of data to be maintained, a level of performance required of the logical volume, or a requirement that one or more snapshots of the logical volume can be preserved representing the data at different points in time. In response to the user's requirements, the system explores the possibilities for providing the characteristics in the existing storage environment. The system then implements the logical volume and stores the intent, along with a corresponding set of rules for configuring the logical volume, with the logical volume. As a result, volume management operations like resizing the volume and evacuating data from the volume use the rules to preserve the intent of the creator. By ensuring that the logical volume consistently conforms to rules fulfilling the original intent, a consistent level of availability of the logical volume can be achieved to fulfill contractual availability requirements for storage service level agreements. (See, e.g., Abstract and paragraph 46.) However, unlike the present invention, Dalal et al., at a minimum, does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 5, the above described third feature of the present invention as recited in independent claim 6, the above described fourth feature of the present invention as recited in independent claim 9, the above described fifth feature of the present invention as

recited in independent claim 10, the above described sixth feature of the present invention as recited in independent claim 11, the above described seventh feature of the present invention as recited in independent claim 16, the above described eighth feature of the present invention as recited in independent claim 19 and the above described ninth feature of the present invention as recited in independent claim 20.

U.S. Patent Publication No. 2004/0193797 (Krehbiel et al.) discloses a method and structure for storing volume and other configuration information on all disk drives of a volume group and for performing pre-merge operations to inform a user of the pending availability of the identified volumes. The invention also discloses a method and system for performing the merge operations and to store all configuration area data in a single data area of the disk drives. More specifically, all configuration information may be stored in the so-called other data area of the configuration information. Insertion of a first disk drive of a volume group may initiate a pre-merge operation by gathering all required information associated with volumes and pseudo volumes of the volume group being inserted. The data so gathered in the pre-merge operation may be presented to the user in advance of insertion of the last drive of a new volume group. Although the pre-merge information may be presented to a user, the user may be precluded from selecting importation of volumes in the newly inserted volume group until all drives of the volume group are inserted into the array storage subsystem. (See, e.g., Abstract and paragraphs 17 - 20.) However, unlike the present invention, Krehbiel et al., at a minimum, does not teach or

suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 5, the above described third feature of the present invention as recited in independent claim 6, the above described fourth feature of the present invention as recited in independent claim 9, the above described fifth feature of the present invention as recited in independent claim 10, the above described sixth feature of the present invention as recited in independent claim 11, the above described seventh feature of the present invention as recited in independent claim 16, the above described eighth feature of the present invention as recited in independent claim 19 and the above described ninth feature of the present invention as recited in independent claim 20.

U.S. Patent Publication No. 2005/0021908 (Ohno et al.) discloses a storage system in which, a storage device such as a disk array subsystem and a host computer are connected via a communication network and a setting method for storage configuration information. Pairs are formed from a plurality of dispersed volumes and copying between the volumes is conducted by a series of remote operations from a management server. A management server instructs the generation of configuration setting files to host computers selected so as to form copy pairs. If the configuration setting files are generated by host agents, the management server instructs copy start to a host computer having a primary volume. The host agent generates a prescribed command via a copy pair control module and causes a storage subsystem to execute copying between the

volumes. When copying is completed, the management server acquires the newest storage configuration information from the storage subsystem and updates configuration information. One embodiment of the invention provides a storage system comprising a host computer, at least one storage device for providing a memory device to the host computer, and a management computer for managing the storage device, where the management computer comprises means for acquiring in advance from the host computer storage configuration information that the host computer has and change notification means for generating change information for changing the storage configuration information that was acquired in advance and posting it to the host computer, and the host computer comprises means for instructing a change of configuration to the storage device based on the change information posted from the management computer. (See, e.g., Abstract and paragraphs 8-15.) However, unlike the present invention, Ohno et al., at a minimum, does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 5, the above described third feature of the present invention as recited in independent claim 6, the above described fourth feature of the present invention as recited in independent claim 9, the above described fifth feature of the present invention as recited in independent claim 10, the above described sixth feature of the present invention as recited in independent claim 11, the above described seventh feature of the present invention as recited in independent claim 16, the above described eighth feature of the present

invention as recited in independent claim 19 and the above described ninth feature of the present invention as recited in independent claim 20.

Therefore, since the cited references at a minimum fail to teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 5, the above described third feature of the present invention as recited in independent claim 6, the above described fourth feature of the present invention as recited in independent claim 9, the above described fifth feature of the present invention as recited in independent claim 10, the above described sixth feature of the present invention as recited in independent claim 11, the above described seventh feature of the present invention as recited in independent claim 16, the above described eighth feature of the present invention as recited in independent claim 19 and the above described ninth feature of the present invention as recited in independent claim 20, and further fail to teach or suggest these features of the present invention in combination with the other limitations recited in each of the specific independent claims, it is submitted that all of the claims are patentable over the cited references whether said references are taken individually or in combination with each other.

In view of the foregoing, Applicant requests that this Petition to Make Special be granted and that the application undergo the accelerated examination procedure set forth in MPEP 708.02 VIII.

Respectfully submitted,

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.



Frederick D. Bailey
Registration No. 42,282

FDB/sdb
(703) 684-1120